

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s):	William F. Micka	Examiner:	Melissa M. Chojnacki
Serial No:	10/079,458	Art Unit:	2164
Filing Date:	February 20, 2002	Docket:	TUC920010091US1(14914)
For:	INCREMENTAL UPDATE CONTROL FOR REMOTE COPY		
Conf. No.:	6646		

DECLARATION PURSUANT TO 37 C.F.R. § 1.131

Sir:

I, William F. Micka, hereby declare that:

1. I am the inventor of the subject matter described and claimed in the above-identified patent application.

2. Prior to March 14, 2001, which is the effective priority date of U.S. Patent No. 6,643,671 to Milillo, I conceived and reduced to practice the invention described in Exhibit A which is the subject of U.S. Patent Application Serial No. 10/079,458.

3. As evidence of the conception and reduction to practice of prior to the date of March 14, 2001, annexed hereto is Exhibit A. Exhibit A is a true photocopy of IBM invention disclosure TUC8-2001-0090 and FCR 572-N, which supports that the invention was created prior to March 14, 2001. The IBM invention disclosure and FCR 572-N contained within Exhibit A includes drawings and an enabling description for the claims of the present application.

All dates within the original IBM invention disclosure have been redacted in the preparation of this Declaration.

4. I do hereby declare that all statements made herein of my own knowledge are true, and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001, and that such willful false statements may jeopardize the validity or enforceability of the patent. The acts relied upon to establish the date prior to the reference or activity were carried out in this country or in a NAFTA country or in a WTO member country.

DECLARATION PURSUANT TO 37 C.F.R. § 1.131

Executed by Inventor 1 of 1

Dec 1, 2009
Dated

W F Micka
William F. Micka

Exhibit A

**Disclosure TUC8-2001-0090**

Prepared for and/or by an IBM Attorney - IBM Confidential

Created By William Micka On
Last Modified By wpts1 wpts1 On
Archived on

Required fields are marked with the asterisk (*) and must be filled in to complete the form.

*** Title of disclosure (in English)**

Incremental Change control for Remote Copy

Summary

Status	Final Decision (File)
Final deadline	
Final deadline reason	
Docket family	TUC8-2001-0091
* Processing location	Tucson
* Functional area	(Myers) Remote Copy Services
Attorney/Patent professional	
Business Area Manager/IDT Lead	
Evaluators	
Submitted date	
* Owning division	SPD
Incentive program	(INC12) Storage Networking
Lab	
* Technology code	
Patent value tool (PVT) score	

Inventors with a Blue Pages entry

Inventors: William Micka/Tucson/IBM

Inventor Name	Serial	Div/Dept	Inventor Phone	Manager Name
Micka, William F.	164898	2D/7XQA	321-4132	

> denotes primary contact

Inventors without a Blue Pages entry**Invention Development Team Information****Main Idea**

To view the Main Idea of this disclosure, open the "Main Idea" document from the view

Critical Questions (Questions 1-9 must be answered in English)**Question 1**On what date was the invention workable? Please format the date as MM/DD/YYYY
(Workable means i.e. when you know that your design will solve the problem)***Question 2**

Is there any planned or actual publication or disclosure of your invention to anyone outside IBM?

☒ Yes
☐ No

If yes, Enter the name of each publication or patent and the date published below.

Publication/Patent: We plan to ship this on Shark, therefore Redbooks and product lit will discuss the function.
 Date Published or Issued: Not published yet.

Are you aware of any publications, products or patents that relate to this invention? ☐ Yes ☒ No

If yes, Enter the name of each publication or patent and the date published below.

Publication/Patent:

Date Published or Issued:

***Question 3**

Has the subject matter of the invention or a product incorporating the invention been sold, used internally in manufacturing, announced for sale, or included in a proposal? ☒ Yes ☐ No

Is a sale, use in manufacturing, product announcement, or proposal planned? ☒ Yes ☐ No

If Yes, Identify the product if known and indicate the date or planned date of sale, announcements, or proposal and to whom the sale, announcement or proposal has been or will be made.

Product: 2105 Shark

Version/Release: Blacktip 2

Code Name:

Date:

To Whom:

If more than one, use cut and paste and append as necessary in the field provided.

***Question 4**

Was the subject matter of your invention or a product incorporating your invention used in public, e.g., outside IBM or in the presence of non-IBMers? ☐ Yes ☒ No

If yes, give a date. Please format the date as MM/DD/YYYY

***Question 5**

Have you ever discussed your invention with others not employed at IBM? ☒ Yes ☐ No

If yes, identify individuals and date discussed. Fill in the text area with the following information, the names of the individuals, the employer, date discussed, under CDA, and CDA #.

Fidelity under CDA. Merrill Lynch under CDA

***Question 6**

Was the invention, in any way, started or developed under a government contract or project? ☐ Yes ☒ No ☐ Not sure

If Yes, enter the contract number

***Question 7**

Was the invention made in the course of any alliance, joint development or other contract activities? ☐ Yes ☒ No ☐ Not Sure

If Yes, enter the following:

Name of Alliance, Contractor or Joint Developer

Contract ID number

Relationship contact name

Relationship contact E-mail

Relationship contact phone

***Question 8**

Have you, or any of the other inventors, submitted this same invention disclosure or similar invention disclosure previously?

☐ Yes

☒ No

If Yes, please provide disclosure number below:

***Question 9**

Are you, or any of the other inventors, aware of any related inventions disclosures submitted by anyone in IBM previously?

☐ Yes

☒ No

If Yes, please provide the docket or disclosure number or any other identifying information below:

Question 10

What type of companies do you expect to compete with inventions of this type? *Check all that apply.*

☒ Manufacturers of enterprise servers

☐ Manufacturers of entry servers

☐ Manufacturers of workstations

☐ Manufacturers of PC's

☐ Non-computer manufacturers

☐ Developers of operating systems

☐ Developers of networking software

☐ Developers of application software

☒ Integrated solution providers

☐ Service providers

☒ Other (Please specify below)

DASD vendors such as EMC and Hitachi

Question 11

If the invention relates to a product or service that is outside the scope of your business unit, please recommend IBM business unit(s), IBM location(s) or individual(s) within IBM that you think would provide a good evaluation of your invention:

***Patent Value Tool (Optional - this may be used by the inventor and attorney to assist with the evaluation)**

(The Patent Value tool can be used by the inventor(s) to determine the potential licensing value of your invention.)

Market

***Question 1:** What is the anticipated annual market size (in dollars) that will be captured by your invention?

Reason(s) for above Answer:

Claims

***Question 1:** How new is the technical field?

Reason(s) for above Answer:

***Question 2:** How central is the invention to the product(s) which might be expected to contain the invention?

Reason(s) for above Answer:

***Question 3:** What is the scope of the claim?

Reason(s) for above Answer:

Portfolio Need

***Question 1:** What are the portfolio needs in the area of your invention?

Reason(s) for above Answer:

Exploitation & Enforcement

***Question 1:** How easily can the use of the invention by a competitor be detected?

Reason(s) for above Answer:

***Question 2:** How easily can the use of the invention be avoided by a competitor?

Reason(s) for above Answer:

Business Value

***Question 1:** What percentage of the companies producing products in the field of this invention might use this invention?

Reason(s) for above Answer:

***Question 2:** What is the value of this patent to current or anticipated Alliance Activity between IBM and other companies?

Reason(s) for above Answer:

***Question 3:** What is the value of this patent to current or anticipated Technology Transfer Activity between IBM and other companies?

Reason(s) for above Answer:

***Question 4:** Does it result in prestige to IBM?

Reason(s) for above Answer:

Evaluation

Search Information

Search Office Information

Allocation and Additional Billers

Final Decision

Post Disclosure Text & Drawings

(Form Disclosure, Revised ())

Main Idea for Disclosure TUC8-2001-0090

Prepared for and/or by an IBM Attorney - IBM Confidential



Archived On

Title of disclosure (in English)

Incremental Change control for Remote Copy

1. Describe your invention, stating the problem solved (if appropriate), and indicating the advantages of using the invention. This invention manages change control for a Remote Data Copy operation. It performs the change control using a combination of FlashCopy and PPRC. Following the initial copy of a data base (or portions thereof), updates to the data are recorded in the FlashCopy structures. These recorded updates are passed to a volume or set of volumes that exist in a PPRC state. Only the data that has changed since the last FlashCopy will be sent to the remote location.

The FlashCopy structure used is shown in the figure as the FlashCopy Bitmap. This bitmap, when the FlashCopy is established with no-background copy, will contain an indication of all updates that were made to the source volume 'A'. The 'B' and 'C' volumes are set up as Peer to Peer copy volumes. Initially, volume(s) 'B' are made equal to volume(s) 'A'. Then PPRC is established between Volume(s) 'B' and 'C' and allowed to complete the entire volume copy. When the copies are complete, volume(s) 'C' are FlashCopied to volume(s) 'C'. In the interval, the volume 'A' bitmaps maintain change recording. At some future time, the customer will want to refresh the Database volume(s) 'D' with updates. This is done by transferring the FlashCopy bitmaps to the PPRC volume(s) 'B' suspension bit maps. The PPRC operation is then told to resynchronize only the data pointed to by bits being set on in the suspension bitmaps. Following the successful copy of the incremental changes, volume(s) 'C' to 'D' can be re-FlashCopied. This can be repeated indefinitely.



AT INCREMENTAL BACKUP TIME

INITIAL SETUP

- FLASH COPY A TO B
 - SELECT NO BKGND COPY
 - FlashCopy Freeze for Consistency
 - When all Freeze, then Thaw all
- PPRC B TO C
 - DETECT FULL DUPLEX STATE
- SUSPEND PPRC B TO C
 - SUSPEND BITMAP MODE
- FlashCopy C TO D
 - SELECT EITHER NO BKGND COPY OR BKGND COPY

- FORCE DESTAGE OF ALL MODIFIED DATA FOR DEV ICE 'A'
- TRANSFER FC BITMAP TO SUSPENDED PPRC PRIMARY
- FC BITMAP INVERTED THEN ORED INTO PPRC SUSPENDED BITMAP
- REESTABLISH FC RELATIONSHIP A TO B
- RESTORE FC PROTECTION BITMAP
- ISSUE INBAND FC COMMAND TO 'B' VOLUME
- RE-ESTABLISH FlashCopy C TO D
 - SELECT EITHER NO BKGND COPY OR BKGND COPY
- PPRC RESYNC B TO C WITH AUTOMATIC SUSPENSION

3. NEW ESTABLISH PPRC Operation

2. How does the invention solve the problem or achieve an advantage, (a description of "the invention", including figures inline as appropriate)?

When a database is copied and maintained current at a remote location, communications lines must be reserved or leased to transfer the data if the data is mirrored on a continuous basis. The time to transfer an entire data base once can be costly in terms of time and communications lines. For instance, to copy a 12 terabyte data base on 4 T3 lines would take around 3 days. The intention of this invention is to provide a mechanism to send only changes since the last backup over the Communication lines following the initial full DB copy. This saves time and provides a more timely copy at the remote location.

3. If the same advantage or problem has been identified by others (inside/outside IBM), how have those others solved it and does your solution differ and why is it better?
This solution creates a true asynchronous remote operation in that there is a total decoupling of data update on the source volumes from transmission of the updates to the remote location. The combination of FlashCopy and PPRC creates a point-in-time copy T0 and the incremental change recording reduces time and cost.

4. If the invention is implemented in a product or prototype, include technical details, purpose, disclosure details to others and the date of that implementation.
Being implemented in Shark for delivery.

FCR 572-N, 'Incremental DSO FlashCopy'

Date Submitted:

Date Approved:

Originator:

Bill Micka

Request Type:

Added new functions

Definition

- Add changes to DSO Establish FlashCopy to specify Incremental Flash Copy operations.
- Add changes to Establish Peer-to-Peer Remote Copy Pair for Incremental FlashCopy in combination with Peer-to-Peer Remote Copy.

Proposed Solution

Change spec as noted below.

See the attachment for functional specification changes.

Persons Contacted Prior to Submitting Request:

Areas Affected by Solution:

Model: All

Devices: All

<input type="radio"/> LOGIC DESIGN	<input checked="" type="radio"/> PROGRAMMING SUPPORT	<input checked="" type="radio"/> PUBLICATIONS
<input checked="" type="radio"/> MICROCODE	<input type="radio"/> DEVICE ENGINEERING	<input type="radio"/> PLANNING
<input checked="" type="radio"/> RAS	<input type="radio"/> DEVICE RAS	<input type="radio"/> FE

Investigator:

W. F. Micka

ADDRESS: 65E/031-1 TUC PHONE: 321-4132

Review Date:

Target Release:

Shark Blacktip

Disposition:

APPROVED

Implementation Responsibility

Shark coders / Host Software / WEB interface

572-N-APPROVED

This FCR contains modifications to the ESS DASD Attachment Specification.

Functional Change

Update Establish Peer-to-Peer Remote Copy Pair order to specify that the Peer-to-Peer Remote Copy operation is in an Incremental FlashCopy - Section 4.11.2.27 Page 132

End of Functional Change

See FCR510

Functional Change

Update DSO Establish FlashCopy order to add the bits and control options for Incremental FlashCopy and FlashCopy in combination with Peer-to-Peer Remote Copy. - Section 4.11.3.3 Page 147

End of Functional Change

4.11.3.3 DSO Establish FlashCopy (Order X'48')

The DSO Establish FlashCopy order identifies to the control unit a source and target volume and the associated extents to create a FlashCopy relationship. There is one type of FlashCopy in the ESS product, Home Area FlashCopy. For Home Area FlashCopy, an entry is created in a source and target relationship table when an extent is established. The entry contains a sequence number. The Sequence number remains the same for all extents established within the current command.

In Phase one and two for Home Area FlashCopy, a volume can have only one relationship at a time. If this command will cause more than one, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 82).

- > If the resources required to complete this command are not available in the control unit or internal hardware conditions prevent the establishment, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 9C).
- > If the number of relationships will exceed the maximum allowed, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 9D). (This is for Data Set level)

In order to ensure the integrity of the FlashCopy relationships, the ESS battery feature must be installed and operational. If it is not, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 83).

The volumes specified in this order must have the same track type. When a full volume FlashCopy copy is specified, the target volume must have greater than or equal to the number of tracks as the source volume. For FlashCopy Phase 1 and 2, the source and target volumes must be in the same logical subsystem. For Phase 3 Home Area FlashCopy only, the number of target extents must equal the number of source extents and the associated tracks must be of the same type. If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 84).

The Nonvolatile Storage must be available. If it is not, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 0D).

For a CKD operation, the addresses in bytes 5 and 7 must be base addresses. If they are not, then execution is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 64).

When the operation is for a CKD volume, the volume specified as the source volume must be the address this command is received on. If the volume specified is a Fixed Block volume, the address this command is received on can be any valid CKD address in the subsystem. If either of these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

Note: If this command is received on a SYS/390 channel and the order indicates that the operation is for a fixed block device, a Diagnostic Control command must have been received previously that enables this capability. See Section 4.9.3.9, "Diagnostic Test Support - F0" on page 89. If the Diagnostic Test Support order Copy Services Mode Control was not set to one, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

- > A Fixed Block LUN specified as a source must be in a ready state. Some reasons for a LUN being not ready are:
 - > • LUN is long busy
 - > • LUN is uninstalled
 - > • LUN is performing a Write Same or Format Unit.
- > If the LUN is not ready, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 65).
- > A Fixed Block LUN specified as a target must not be reserved to another initiator. If the LUN is reserved to another initiator, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 9B).
- > A Fixed Block LUN specified as a target must not be long busy, uninstalled, or performing a Write Same or

- > Format Unit. If any of these conditions exist, the
- > command is terminated with status that includes unit
- > check (Command Reject, format X'0F', Status Not as
- > Required, Reason Code 13).

The volume specified in bytes 4-5 or 6-7 must not be in CE mode, CC3 (non-existent device), write inhibited by internal Controller action, device long busy, or 'Status cannot be determined' state. If they are, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 22).

If the specified target volume is a Peer-to-Peer Remote Copy primary or secondary volume or a XRC source volume, and bytes 2-3, bit 10 is zero, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 85).

If the specified source or target volume has pinned data, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 56).

When a Diagnostic Control Suborder 15 has been received by a device, the device is protected from being overwritten by outboard copy functions. If a device in this state is specified as a FlashCopy target and the FlashCopy feature has been enabled, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 87). Receipt of a Diagnostic Control Suborder 16 will clear the protected state.

- > If byte 1 bit 1 is set indicating a FlashCopy Remote
- > operation:
 - > • The volume receiving the command must be a
 - > Peer-to-Peer Remote Copy primary. If this condi-
 - > tion is not met, the command is terminated with
 - > status that includes unit check, (Command Reject,
 - > format X'0F', Status Not as Required, Reason
 - > Code 94).
 - > • The Secondary Peer-to-Peer Remote Copy volume
 - > must not be long busy. If it is, the command is ter-
 - > minated with status that includes unit check
 - > (Command Reject, format X'0F', Status Not as
 - > Required, Reason Code 13).
 - > • The FlashCopy Source volume address in bytes
 - > 4-5 must be the same as the secondary PPRC
 - > volume associated with the PPRC primary volume
 - > receiving this command. If this condition is not
 - > met, the command is terminated with status that
 - > includes unit check, (Command Reject, format
 - > X'0F', Status Not as Required, Reason Code A0).
 - > • The Secondary Peer-to-Peer Remote Copy volume
 - > must not be long busy. If it is, the command is ter-
 - > minated with status that includes unit check
 - > (Command Reject, format X'0F', Status Not as
 - > Required, Reason Code 13).
 - > • The secondary Peer-to-Peer Remote Copy volume
 - > must not be in 'Status cannot be Determined'
 - > state. This information will be obtained by a
 - > RSSD order X'0F' command sequence from the
 - > primary CU to the Secondary CU over the Peer to
 - > peer interface. If it is, the command is terminated
 - > with status that includes unit check (Command

- > Reject, format X'0F', Status Not as Required,
- > Reason Code 22).

- > • There must be at least one active path estab-
- > lished to the secondary control unit to complete
- > this order. If there are no paths, the command is
- > terminated with status that includes unit check
- > (Command Reject, format X'0F', Status Not as
- > Required, Reason Code 52).
- > • If the primary control unit timed out waiting for
- > communications resources to complete the
- > command with the secondary volume, the
- > command is terminated with status that includes
- > unit check (Command Reject, format X'0F', Status
- > Not as Required, Reason Code 5A).

If byte 1 bit 3 is set indicating a Restart FlashCopy operation, the volumes specified as the source and target must be the same as an existing relationship for either volume. If the specified source and target are not in any FlashCopy relationship, this command establishes the FlashCopy as if byte 1 bit 3 was not set. If the source and/or target volumes are involved in another relationship, the command is terminated with status that includes unit check, (Command Reject, format X'0F', Status Not as Required, Reason Code 97).

- > If byte 1 bit 3 is set indicating a Restart FlashCopy
- > operation and the operation could not be completed
- > due to the inability to write metadata or global status,
- > the operation will be terminated, the volumes will be
- > placed into a terminate pending state with an indi-
- > cation in NVS to this effect. When the fault causing
- > this error is repaired, the volumes will be terminated.
- > In the terminate pending state, status commands will
- > sense that the volumes are simplex and new
- > FlashCopy relationships can be established. When the
- > terminate pending state is entered, the command is
- > terminated with status that includes unit check,
- > (Command Reject, format X'0F', Status Not as
- > Required, Reason Code 9A).
- > If bytes 2-3 bit 10 is set indicating an Incremental
- > Peer-to-Peer Remote Copy operation, the target
- > volume must be a Suspended Peer-to-Peer Remote
- > Copy Primary volume. The Primary volume must have
- > received a Establish Peer-to-Peer Remote Copy Pair
- > order with bytes 32-33 bit 7 set to one. If these condi-
- > tions are not met, the command is terminated with
- > status that includes unit check, (Command Reject,
- > format X'0F', Status Not as Required, Reason Code
- > 98).

A CKD target device must not be in an on-line state while a DSO Establish FlashCopy is attempting to establish the pair when byte 2 bit 5 is set to 1. If it is, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 60).

For the specified CKD volume, when this command is first-in-chain (excluding SMR), and a Perform Subsystem Function, Set Special Intercept Condition was executed for the addressed device on any interface with the same path group identifier as the interface processing the command, processing of this command is terminated with status that includes unit check (Command Reject, Environmental Data Present,

format X'0F' Reason Code 80) with a program action code of X'70'.

This order requires a minimum of twenty six parameter bytes. The maximum number of bytes that can be sent is 1786. The order-unique bytes are defined as follows:

Figure 20. DSO Establish FlashCopy Parameters		
Byte	Value	Description
0	X'48'	DSO Establish FlashCopy
1		FlashCopy Flags Byte 1
		Bits Description
		0 Message required - set to one for FlashCopy Phase 1 & 2. A message will be provided when the 'destage of modified data' task from the source volume's cache has completed and the 'discard of modified data' task for the target has been completed.
		1 FlashCopy Remote. See detail following the figure.
		2 FlashCopy Freeze. See detail following the figure.
		3 FlashCopy Restart. See detail following the figure.
		4-7 Not used
2-3		FlashCopy Flags Bytes 2-3. See text following table for more detail.
		Bits Description
		0 Not used
		1 Accelerated destage mode.
		2 Peer-to-Peer Remote Copy mode.
		3 CKD volume if set to zero. Fixed Block volume if set to one.
		4 If zero, perform a background copy for the identified extents. If set to one, do not perform the background copy.
		5 When this bit is set to zero, allow the target device to be in an on-line state. When this bit is set to one, do not allow the pair to be established if the target volume is on-line to any host.
		6 If this bit is set to zero, destage all source cache modified data before allowing modifications to the volume.
		7 Start Change Recording for the FlashCopy volume pair.
		8 Inhibit writes to the target volume.
		9 Persistent FlashCopy Relationship.
		10 Incremental Peer-to-Peer Remote Copy operation.
		11 Perform operations at a specific time interval when set to one.
		12 Restore FlashCopy pair with changed tracks indicated in the change recording bitmap. (Not a BlackTip deliverable.)
		13-15 Not used
4-5		Source Volume address. High byte is the LSS. Low byte is the CKD volume address or the Fixed Block LUN address. When byte 1 bit 1 is set to one, the address specified is that of a remote Peer-to-Peer Remote Copy secondary volume.
6-7		Target Volume address. High byte is the LSS. Low byte is the CKD volume address or the Fixed Block LUN address. When byte 1 bit 1 is set to one, the address specified is that of a remote volume that will be the target volume for a Peer-to-Peer Remote Copy secondary.
8-9		Peer-to-Peer Remote Copy target/secondary volume SSID. This field is valid only when byte 1, bit 1 or FlashCopy Flags bytes 2-3 bit 2 is set to one.
10-13		Peer-to-Peer Remote Copy target/secondary volume control unit serial number. This field is valid only when byte 1, bit 1 or FlashCopy Flags bytes 2-3 bit 2 is set to one.

Figure 20. DSO Establish FlashCopy Parameters												
Byte	Value	Description										
20-21		Source host device number. Used by VM for mini disk identification. This field is not used or stored in the subsystem.										
22-23		Target host device number. Used by VM for mini disk identification. This field is not used or stored in the subsystem.										
24-nn		These fields have a different meaning depending on the Flag byte and the FlashCopy flags. If byte 1, bit 1 or FlashCopy Flags bits 10 or 11 are set to one, the count of extents and the extents fields will be shifted down in order to provide more parameter bytes.										
24-25		Format when byte 1 bit 1 and FlashCopy flags bits 10 and 11 are all zero. For a CKD volume, a count of the Extent pairs to follow. The maximum value is 110. For a Fixed Block LUN, this field is not used. A value of zero is allowed and performs parameter checking before CE/DE but is a NOP for all other functions.										
26-41		Extent Pairs in the form CCHH. <table><tr><th>Bytes</th><th>Description</th></tr><tr><td>26-29</td><td>Source Start.</td></tr><tr><td>30-33</td><td>Source End.</td></tr><tr><td>34-37</td><td>Target Start.</td></tr><tr><td>38-41</td><td>Target End.</td></tr></table>	Bytes	Description	26-29	Source Start.	30-33	Source End.	34-37	Target Start.	38-41	Target End.
Bytes	Description											
26-29	Source Start.											
30-33	Source End.											
34-37	Target Start.											
38-41	Target End.											
42-nn		Remaining Extent pairs.										
24-41		Format when byte 1 bit 1 or FlashCopy flags bits 10 or 11 are set to one. This is not a SlackTip deliverable.										
24-25		FlashCopy Flags Extended. <table><tr><th>Bit</th><th>Description</th></tr><tr><td>0</td><td>The time field is valid and contains an 8 byte clock value to be compared with the ESS internal clock. When the values are equal or the ESS clock is greater than the value, a FlashCopy operation will be performed.</td></tr><tr><td>1</td><td>The time field is valid and the first 4 bytes contain a count. The granularity of each count is 100 milliseconds. When the ESS achieves a value which is a multiple of the number in this field multiplied by 100 milliseconds, a FlashCopy operation will be performed.</td></tr><tr><td>2-16</td><td>Not used</td></tr></table>	Bit	Description	0	The time field is valid and contains an 8 byte clock value to be compared with the ESS internal clock. When the values are equal or the ESS clock is greater than the value, a FlashCopy operation will be performed.	1	The time field is valid and the first 4 bytes contain a count. The granularity of each count is 100 milliseconds. When the ESS achieves a value which is a multiple of the number in this field multiplied by 100 milliseconds, a FlashCopy operation will be performed.	2-16	Not used		
Bit	Description											
0	The time field is valid and contains an 8 byte clock value to be compared with the ESS internal clock. When the values are equal or the ESS clock is greater than the value, a FlashCopy operation will be performed.											
1	The time field is valid and the first 4 bytes contain a count. The granularity of each count is 100 milliseconds. When the ESS achieves a value which is a multiple of the number in this field multiplied by 100 milliseconds, a FlashCopy operation will be performed.											
2-16	Not used											
26-27		Session ID when not zero.										
28-35		Time value to used for a scheduled FlashCopy operation.										
42-nn		Not used when byte 1 bit 1 or FlashCopy flags bits 10 or 11 are set to one. These fields could be used when Data Set FlashCopy is designed.										

Note: A specific track or extent is addressed by a two-level address of the form 'CCHH', where 'CC' is an unsigned 16-bit binary cylinder number and 'HH' is an unsigned 16-bit binary track number within the cylinder.

When a Fixed Block volume is specified, bytes 24 and following are ignored if received for the command.

Detail for Flag byte 1

Bits Description

0 Message required.

Note: In addition, this bit will be set when FlashCopy Remote is specified and when the command is issued to the Peer-to-Peer Remote Copy secondary volume. The Peer-to-Peer Remote Copy interface must provide the read message ID function.

1 FlashCopy Remote. With this bit set to one, the source and target volumes are referenced to a Peer-to-Peer Remote Copy secondary location. This command must be issued to a Peer-to-

Peer Remote Copy primary volume when this bit is set. Full volume only operations are allowed.

When this command is received with this bit set to one, the Peer-to-Peer Remote Copy pair will become suspended following communication of the FlashCopy order to the remote location.

The remote volume is considered a conduit to transmit the command. The actual volume address is contained in the parameters.

- 2 FlashCopy Freeze. Following the destage of source data (during which the volume is not accepting writes), the volume will remain long busy until the thaw command (Consistency Group Created) releases the volume or the timer expires.

The timer is set to a default of 120 seconds and is called the Consistency Group timer on the WEB panel that displays server properties. The customer or delegate can change the time value to their choice. The timer is used for both Peer-to-Peer Remote Copy and FlashCopy. It is activated when a Peer-to-Peer Remote Copy volume is suspended and when FlashCopy Freeze option is chosen and Consistency Grouping has been activated for the logical subsystem. The timer will be stopped when a Consistency Group Created order is received.

- 3 FlashCopy Restart. With this bit set to one, the FlashCopy relationship can be restarted. Only one relationship is allowed on a volume and this bit will allow a refresh of the data without requiring a DSO Withdraw FlashCopy order to be executed. The volumes specified must be the same as in the previous command execution. This bit is ignored if there is no FlashCopy relationship on the volumes.

If this command originates from the Copy Services WEB server, only full volume FlashCopy will be specified.

Detail for Flag bytes 2-3

Bits Description

- 0 Not used
- 1 Accelerated destage mode. This bit being set will allow data staged into cache for a Concurrent Copy operation (for example) to be demoted sooner than a normal LRU destage. A property of a dump or copy operation is that the data will only be read once.
- 2 Peer-to-Peer Remote Copy mode. In Phase one, this bit must be set to zero.
- 3 CKD volume if set to zero. Fixed Block volume if set to one.

Note: If this command is received on a SYS/390 channel and the order indicates that the operation is for a fixed block device, a Diagnostic Control command must have been received previously that enables this capability. See Section 4.9.3.9, "Diagnostic Test Support - F0" on page 89.

- 4 If zero, perform a background copy for the identified extents. If set to one, do not perform the background copy.

- 5 When this bit is set to zero, allow the target device to be in an on-line state.

When this bit is set to one, do not allow the pair to be established if the target volume is on-line to any host.

Note: It is assumed that this is detected by the absence of any path group state which indicates the device is grouped. This bit is ignored if a Fixed Block device is indicated.

- 6 If this bit is set to zero, destage all source cache modified data before allowing modifications to the volume. In Phase one and two, this bit must be set to zero.

- 7 Start/Continue Change Recording for the FlashCopy volume pair when set to one. Change recording for the entire volume will begin on both the source and target volumes. If Flag byte bit 8 (Inhibit write of the target) is set to one, change recording will not be performed on the target volume.

Note: For Blacktip, the implementation will perform change recording using the FlashCopy bitmap which is only on the target volume. The bitmap will be valid for change recording only following a destage of all source cache modified data to the physical volume.

When this bit is zero and the pair is performing change recording,

if bit 12 is set the bitmaps will toggle and change recording will stop.

if bit 12 is zero the bitmaps will be unchanged and change recording will stop.

- 8 Inhibit writes to the target volume. The sense given will be the same as if the volume were a Peer-to-Peer Remote Copy secondary.

- 9 Persistent FlashCopy Relationship. When set to one, the relationship can only be removed by a DSO Withdraw FlashCopy order.

- 10 Incremental Peer-to-Peer Remote Copy operation. For this operation, the Peer-to-Peer Remote Copy primary volume is the target of a FlashCopy. When this bit is set to one, the 'n-1' change recording bitmap will be 'ored' with the Peer-to-Peer Remote Copy Suspension bit map. Incremental operations operate only on full volumes.

To perform the actual resynchronization, a Establish Peer-to-Peer Remote Copy Pair order with flag byte 1 set to copy out of sync tracks to the secondary must be executed.

- 11 Perform operations at a specific time interval when set to one.

This is not a BlackTip deliverable.

- 12 Restore FlashCopy pair with changed tracks indicated in the change recording bitmap. The source and target change recording bitmaps will be 'ored' and the incremental FlashCopy

will restore from the source tracks. (**Not a BlackTip deliverable.**)

Note: The notion of restoring from either the source or target volume can be accomplished by correctly specifying the volume in the source and target parameters of this command.

13-15 Not used

Sparse Volume: The extents in bytes 26-41 (and following) contain track addresses to identify one or more tracks to copy or make available to the target volume. Phase one of FlashCopy performs operations on a single volume to single volume relationship but not all tracks need to be specified on the volumes. Those tracks not specified will not be copied to the target and a read or non-format write to these non-specified tracks will provide unpredictable results.

Note: The DSO Establish FlashCopy order will cause cache modified data for the identified extents to be destaged to the source volume and discarded for target volume. Data for tracks not identified will be maintained as if FlashCopy was not executing.

The CCHH values in the start extents and end extents parameters must be valid track addresses for the volume and the end extent must be greater than or equal to the matching start extent. See Appendix A, "Device data, geometries, and facilities" on page 313. If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

For Phase 1 of FlashCopy, byte 2 bits 2 and 6 must be zero and the target start and end extents must be the same as the source start and end extents. For Phase 1 & 2, byte one bit zero must be set to one. If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

Only a Read Message ID can be chained if the Message Required flag is one. If this condition is not met, the command is terminated with status that includes unit check (Command Reject, format X'02', Invalid Command Sequence).

When the Restore FlashCopy pair order is specified (bytes 2-3 bit 12) and the operation is to restore from the original target back to the source, certain restrictions are required. (**Not a BlackTip deliverable.**)

- A FlashCopy relationship must exist for the two volumes.
- The FlashCopy relationship must be persistent.
- The background copy must be specified and completed.

If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 8A).

Channel End and Device End are returned when:

1. the parameters have been verified

2. the source 'destage cache modified data' task has been queued
3. the target 'discard cache modified data' task has been queued

Commands that attempt to modify either the source or target volume prior to completion of the destage/discard task will receive a long busy unit check. Long Busy will be given until destage is complete.

Following device end presentation, the asynchronous task of destaging data will proceed. If that operation ended abnormally due to pinned data, the message presented will be X'06'.

For Phase 3, long busy will not be used since the destage of source cache modified data will not be completed prior to allowing write access to the volume. This will improve the time to complete the command. If a write occurs to a track that has modified data not yet written to the device, CCR will be signalled on the channel and an immediate destage will be performed for the affected track.

Extended Long Busy State: The Long Busy (State Change Pending) state is designed to keep an I/O for a volume queued in the host until the controller finishes with work in progress that would normally take longer than a missing interrupt timer interval. When DSO Establish FlashCopy is received, all the volume(s) that are affected will destage the source cache modified data. The volume(s) will be in the long busy state until that destage work is completed. For cases where the customer wants to create a consistent FlashCopy across many volumes in the same or different controllers, a capability is provided to Freeze the volumes and not allow updates. Bit 2 in Flag byte 1 will cause the devices to go into a timed (WEB panel consistency grouping timer) Extended Long Busy state following the destage activity. The reason for this is to allow enough time for a controlling host to send DSO Establish FlashCopy order with Freeze indicated to all controllers followed by Consistency Group Created to all controllers. Receipt of Consistency Group Created order will stop the timer and reset long busy. The Freeze operation can be issued without regard to activity happening on the volumes. The idea is to stop dependent I/O from executing on volumes that have already been FlashCopied until the dependent volume has been FlashCopied. Once all volumes of interest have received this command, I/O can be allowed and the target volumes are all consistent at a point in time.

Since the busy state affects customer applications, the Freeze and subsequent Thaw should execute very rapidly. Each command should complete in less than a half second. The time value for the Extended Long Busy is set by default to 2 minutes. A Consistency Group Created will reset the Extended Long Busy state if being actively timed.

Note: The Extended Long Busy state is called 'Consistency Grouping' on the WEB Copy Services Screens. The Extended Long Busy state being set on will only affect the initiation of an additional timing interval following the normal destage activity for an individual volume. DSO Establish FlashCopy and Consistency Group

Created will be accepted when the Extended Long Busy state is not set.

Programming Note: When this order is issued under VM, the target volume is protected from user specification error by checking to see if the issuer of the order has write access to the volume. This is done using the values in bytes 20-23 which are set by the MVS guest to be the device address it knows.

Note: Prior to completion of this order, all modified data in cache for the specified source volume extents must be destaged and all modified or unmodified data in cache for the specified target volume extents must be discarded. In the extreme case, where the NVS is full of small records and the customer has started many establish orders, the time to destage can be quite long and could exceed the missing interrupt handler time out. For this reason, the command must execute in an asynchronous manner. The message required bit in the flag byte provides this capability. In Phase 3, the intention is to create another structure in the controller that will allow a deferred destage thus gaining better performance on the command.

Functional Change

Update DSO Withdraw FlashCopy order to add the bits and control options for allowing this command to be sent in a Peer-to-Peer Remote Copy relationship. - Section 4.11.3.4 Page 149

End of Functional Change

4.11.3.4 DSO Withdraw FlashCopy (Order X'49')

The DSO Withdraw FlashCopy order identifies to the control unit a source and target volume and the associated extents to withdraw from a FlashCopy relationship.

Note: This order performs no operation on a LSA volume pair when operating in native LSA mode unless the volume has been placed in bit map mode. CE/DE status is returned upon determination that the volumes are in a native LSA subsystem and are not operating in bit map mode..

For a CKD operation, the addresses in bytes 5 and 7, when specified, must be base addresses. If they are not, then execution is terminated with status that

includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 64).

Address Checking

- When the operation is for a CKD volume and byte 2 bit 2 is zero, the volume specified as the source volume must be the address this command is received on.
- If byte 2 bit 2 is set to one, the volume specified as the target in bytes 6-7 must be the volume this command is received on.
- If the volume specified is a Fixed Block volume, the address this command is received on can be any valid CKD address in the subsystem.

If any of these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

Note: If this command is received on a SYS/390 channel and the order indicates that the operation is for a fixed block device, a Diagnostic Control command must have been received previously that enables this capability. See Section 4.9.3.9, "Diagnostic Test Support - F0" on page 89. If the Diagnostic Test Support order Copy Services Mode Control was not set to one, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

For the specified CKD volume, when this command is first-in-chain (excluding SMR), and a Perform Subsystem Function, Set Special Intercept Condition was executed for the addressed device on any interface with the same path group identifier as the interface processing the command, processing of this command is terminated with status that includes unit check (Command Reject, Environmental Data Present, format X'0F' Reason Code 80) with a program action code of X'70'.

This order requires a minimum of twenty six parameter bytes. The maximum number of bytes that can be sent is 1786. The order-unique bytes are defined as follows:

Figure 21 (Page 1 of 2). DSO Withdraw FlashCopy Parameters		
Byte	Value	Description
0	X'49'	DSO Withdraw FlashCopy
1		Flags
	Bit	Description
	0	Not used
	1	FlashCopy Remote. This command will be sent to the Peer-to-Peer Remote Copy secondary.
	2-7	Not used

Figure 21 (Page 2 of 2). DSO Withdraw FlashCopy Parameters

Byte	Value	Description
2-3		FlashCopy Flags
	Bits	Description
	0	When set to one, withdraw specified extents. Note: Bit 0 will be set to zero for FlashCopy Phase 1.
	1	When set to one, withdraw all extents on the source volume and all extents on the target volume. This function is provided to remove all extents in the case when knowledge of which extents are active has been lost or is not available.
	2	Specify target volume only. With this bit set to one, only the target address is required to be specified in the parameters. The source volume address in bytes 4-5 will be determined within the control unit from stored information.
	3	CKD volume if set to zero. Fixed Block volume if set to one. Note: If this command is received on a SYS390 channel and the order indicates that the operation is for a fixed block device, a Diagnostic Control command must have been received previously that enables this capability. See Section 4.9.3.9, "Diagnostic Test Support - FC" on page 89.
	4	Stop Change Recording if set to one. A value of zero will leave change recording if it is active on either volume. (Not a BlackTip deliverable.)
	5-15	Not used
4-5		Source Volume address. High byte is the LSS. Low byte is the CKD volume address or the Fixed Block LUN address. Note: When byte 2 bit 2 is set to one, this field should be set to zero by the invoker of the command. The control unit will not check the values for zero.
6-7		Target Volume address. High byte is the LSS. Low byte is the CKD volume address or the Fixed Block LUN address.
8-9		Peer-to-Peer Remote Copy target volume SSID. Valid in Phase 2. This field is valid only when byte 1 bit 1 is set to one.
10-19		Peer-to-Peer Remote Copy target volume control unit serial number. Valid in Phase 2. This field is valid only when byte 1 bit 1 is set to one.
20-21		Source host device number. Used by VM for mini disk identification. This field is not used or stored in the subsystem.
22-23		Target host device number. Used by VM for mini disk identification. This field is not used or stored in the subsystem.
24-25		Count of Extent Pairs. For Phase one, this value will be zero. Byte 2 bit 1 will be set to one to indicate full volume withdraw only. For Phase two and following phases, byte 2 bit zero would be valid and would indicate that the following extent ranges would be used. The maximum value is 110.
26-41		Extent Pairs. CCHH or LBA track descriptor.
	Bytes	Description
	26-29	Source Start.
	30-33	Source End.
	34-37	Target Start.
	38-41	Target End.
42-4n		Remaining Extent pairs.

When used, the CCHH values in the start extents and end extents parameters must be valid track addresses for the volume type and the end extent must be greater than or equal to the matching start extent. See Appendix A, "Device data, geometries, and facilities" on page 313. If these conditions are

not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

- > If byte 1 bit 1 is set indicating a Remote Peer-to-Peer Remote Copy operation:
 - > • The volume receiving the command must be a Peer-to-Peer Remote Copy primary. If this condition is not met, the command is terminated with status that includes unit check, (Command Reject, format X'0F', Status Not as Required, Reason Code 94).
 - > • The Secondary Peer-to-Peer Remote Copy volume must not be long busy. If it is, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 13).
 - > • The secondary Peer-to-Peer Remote Copy volume must not be in 'Status cannot be Determined' state. This information will be obtained by a RSSD order X'0F' command sequence from the primary CU to the Secondary CU over the Peer to peer interface. If it is, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 22).
 - > • There must be at least one active path established to the secondary control unit to complete this order. If there are no paths, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 52).
 - > • If the primary control unit timed out waiting for communications resources to complete the command with the secondary volume, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 5A).

For Phase 1 and 2 of Home Area FlashCopy, byte 2 bit 0 and bytes 24-25 must be zero; byte 2 bit 1 must be one. If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

If a relationship exists for the source device and byte 2 bit 2 is set to zero, the values in bytes 4-19 must be the same as received in a prior DSO Establish FlashCopy order. If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

If a relationship exists for the target device and byte 2 bit 2 is set to one, the source volume will be determined from Copy Services global status and this command will cause the removal of the relationship between the source and target volumes.

Note: When data set level FlashCopy is implemented, the extents specified may be used to withdraw portions of the relationships.

If a relationship exists for the target device with byte 2 bit 2 set to one, and the source volume cannot be determined from Copy Services global status, the command will remove the relationship for the target device and good ending status will be returned.

If there is no relationship for the specified operation, the command will perform no operation and good ending status will be returned.

Channel End is returned when the parameters have been verified. Device End is returned when the extents have been withdrawn and the relationship removed.

Functional Change

Change Consistency Group Created order to make it work for both PPRC and FlashCopy - Section 4.11.2.34 Page 141

End of Functional Change

4.11.2.34 Consistency Group Created (Order X'67')

This order is designed for Data Base Consistency Automation. The receipt of this order signals to the controller that software has completed the work necessary to ensure secondary PPRC volume consistency and FlashCopy volume consistency across all controllers and devices that are related.

For Peer-to-Peer Remote Copy volumes. Within an individual primary control unit, devices are related if they have the same secondary control unit which is identified by its sequence number and SSID.

For FlashCopy volumes.: This command will apply to all the FlashCopy volumes in the logical subsystem or in the session when Flag byte bit 1 is set to one. Bytes 16-29 are not used in this mode.

For both type of volumes.: Those volumes that were in an Extended Long Busy state will clear that state and begin to accept I/O activity again. Those same volumes will clear the *potential* to go Extended Long Busy so that a subsequent update I/O that would have triggered this state will execute without application interruption.

This order can be issued on any installed volume address (should be a Utility Device which cannot be reserved to another host).

This order will be executed even if the volume address used on the channel is in the long busy state.

If a "PSF or DSO Set Special Intercept Condition" was executed for the addressed device on any interface with the same path group identifier as the interface processing the command, processing of this command is terminated with status that includes unit check (Command Reject, Environmental Data Present, Format X'0F', Reason Code X'80') with a program action code of X'70'. See 4.11.2.6, "Set Special Intercept Condition (Order X'1B') on page 111.

This order requires 32 parameter bytes to be transferred. The order unique-bytes are defined as follows:

Figure 22. Command Format - Consistency Group Created												
Bytes	Value	Description										
0	X'67'	Consistency Group Created										
1	Flags	<table><tr><th>Bits</th><th>Description</th></tr><tr><td>0</td><td>Message required - set to zero.</td></tr><tr><td>1</td><td>Execute this command for FlashCopy volumes when set to one. All volumes in the logical subsystem or session will thaw.</td></tr><tr><td>2</td><td>Session ID in bytes 30-31 valid if set to one. (Not a BlackTip deliverable.)</td></tr><tr><td>3-7</td><td>Not used</td></tr></table>	Bits	Description	0	Message required - set to zero.	1	Execute this command for FlashCopy volumes when set to one. All volumes in the logical subsystem or session will thaw.	2	Session ID in bytes 30-31 valid if set to one. (Not a BlackTip deliverable.)	3-7	Not used
Bits	Description											
0	Message required - set to zero.											
1	Execute this command for FlashCopy volumes when set to one. All volumes in the logical subsystem or session will thaw.											
2	Session ID in bytes 30-31 valid if set to one. (Not a BlackTip deliverable.)											
3-7	Not used											
2		Primary Logical Subsystem number										
3		Not used										
4-13		Primary control unit Sequence number										
14-15		Primary control unit SSID										
16		Secondary Logical Subsystem number										
17		Not used										
18-27		Secondary Control Unit Sequence number										
28-29		Secondary control unit SSID										
30-31		Session ID (Not a BlackTip deliverable.)										

For a CKD device, the LSS value in byte 2 must be the same as the LSS the command was received on. If this condition is not met, the command is terminated with status that includes unit check (Command Reject, format X'04', Invalid Parameter).

Byte 2 must contain the correct Primary logical subsystem number; bytes 4-13 must contain the correct Primary CU sequence number; bytes 14-15 must contain the correct Primary CU SSID. For Peer-to-Peer Remote Copy only, byte 16 must contain the correct Secondary logical subsystem number; bytes 18-27 must contain the correct Secondary CU sequence number; bytes 28-29 must contain the correct Secondary CU SSID; If these conditions are not met, the command is terminated with status that includes unit check (Command Reject, format X'0F', Status Not as Required, Reason Code 53).

For FlashCopy, bytes 16-29 are not used.

The Session ID, if valid, will be used to identify the volumes to be affected by this command.

Channel End and Device End will be given when the parameters have been validated.

Functional Change

Add a bits to the RSSD FlashCopy Relationship Table for the device to indicate persistent copy and long busy.- Section 4.11.5.18 Page 171

End of Functional Change

4.11.5.18 FlashCopy Relationship Table for the Device

If byte 6 of the Perform Subsystem Function Prepare for Read Subsystem Data is '11', the entire FlashCopy Relationship Table for the volume will be read. Each relationship is 26 bytes and when 1024 relationships (maximum) are established, this command will transfer 26,628 bytes. A minimum of 4 bytes will be transferred. When bytes 0-1 contain zero, the information in bytes 2-3 is valid and bytes 4-29 will be unused and set to zero.

Note: When a FlashCopy has been established between two LSA volumes in the same logical subsystem and Incremental Copy has not been specified, the Relationship Table will be empty and so indicated by a zero value in bytes 0-1.

Figure 23. FlashCopy Source or Target Relationship Table

Byte	Value	Description
0-1		Number of Entries for this volume. Set to zero when there are no FlashCopy relationships active. For Phase One of FlashCopy, the maximum value is one.
2		Size of each entry. Set to X'1A'.
3		FlashCopy Version
	Bits	Description
	0-4	Not used
	6	Volume is long busy for either the destage of source data or due to an extended long busy for consistency grouping.
	5	Volume is extended long busy for consistency grouping.
	6-7	
	00	Phase one FlashCopy only supported.
	01	Phase two FlashCopy supported.
	10	Phase three FlashCopy supported.
4-29	These bytes constitute a relationship entry.	
4-5	LSS / Volumes	Source or Target Address
6		Flags
	Bits	Description
	0	When zero, this is a source relationship table. When set to one, it is a target relationship table.
	1	When zero, no Peer-to-Peer Remote Copy link required. When set to one, a Peer-to-Peer Remote Copy link is required for the volume pair.
	2	When zero, background copy task for the relationship is performed. When set to one, no background copy will be done.
	3	When set to one, a background copy task is actively being performed. Note: This bit is needed since background throttling came to be.
	4	Extents are in a Persistent FlashCopy Relationship
	5	Target Extents are Write Protected
	6-7	Not used
7-11		Sequence number or time value
12-13	SSID	Target Logical Subsystem SSID value.
14-17	Count	Number of contiguous tracks in the extent or for Phase 1 see note below.
18-21	CCHH	Source Start Extent
26-29	Count	Number of tracks remaining to copy for the extent.
	The format of bytes 4-29 is repeated the number of times indicated in bytes 0-1.	

Note: For Phase 1, bytes 14-19 contain a value which is equal to the number of tracks on the full volume.

Functional Change

Add a bit to the Sense command to indicate the sense is for the Remote volume, not the local one.-
Section 6.3.1.1 Page 201

End of Functional Change

The sense bytes are defined as follows:

Figure 24. Sense Information summary

Byte	Bit	Meaning	Page
0	0	Command Reject	199
	1	Intervention Required	
	2	Bus Out Parity Check	
	3	Equipment Check	
	4	Data Check	
	5	Overrun	
	6	Not used	
	7	Incomplete Domain	
1	0	Permanent Error	199
	1	Invalid Track Format	
	2	End of Cylinder	
	3	Message to Operator	
	4	No Record Found	
	5	File Protected	
	6	Write Inhibited	
	7	Imprecise Ending	
2	0	Request Inhibit Write	199
	1	Correctable (data check)	
	2	First Error Log	
	3	Environmental Data Present	
	4	Not used	
	5	Imprecise Ending	
	6	Not used	
	7	Sense information for Remote Volume	
3		Residual Count	201
4		Physical Drive ID	201
5		Low Cylinder Address	201
6		Head Address and High Cylinder Address	201
7	0-3	Format Message	202
	4-7		
8-23		Format Dependent Information	203 - 209
24		Only used in ECKD 32-Byte format	202
25		Program Action Code	202
26-27		Configuration Information	202
28		Message Code	203
29-31		Cylinder and Head Address	203

Functional Change

Add description of byte 2 bit 7 - Section 6.3.1.1
Page 201

End of Functional Change

Byte 2, Bit 7 Sense for Remote Volume: When this bit is set to one, the sense information returned is for a volume that is addressed at a Peer-to-Peer Remote Copy secondary control unit by using a Peer-to-Peer Remote Copy pair and its communications link to send the command.

Functional Change

Add additional reason codes for format 0 message F
- Section 6.3.2.1 Page 203

End of Functional Change

- > 13 An attempt to establish a CKD PPRC pair and the secondary device is busy, reserved, or participating in a Concurrent Copy or XRC session. When attempting to establish a Fixed Block PPRC pair or FlashCopy pair, the LUN cannot be long busy, uninstalled, waiting for a Start Unit command, with Format Failed status, active on another interface, in a hardware failure state, or performing a Write Same or Format Unit.
- 80 Attempt to execute a Global Command on an interface that was disabled for Global Commands by a Perform Subsystem Function, Set Special Intercept Condition order.
- 81 An attention was presented on an interface that was disabled for Global Commands by a Perform Subsystem Function, Set Special Intercept Condition order.
- 82 No resources available to create relationships for the DSO Establish FlashCopy command or the number of relationships will exceed the maximum allowed.
- 83 ESS battery feature unavailable or not installed. FlashCopy requires the battery for proper operation. DSO Establish FlashCopy command.
- 84 The FlashCopy volumes specified are not in the same Logical Subsystem or they are not of the same size and format.
- 85 The target FlashCopy volume is a Peer-to-Peer Remote Copy primary or secondary or is an XRC primary volume.
- 86 The feature that the command references is not installed or enabled for the subsystem.
- 87 Outboard copy operations are not allowed to overwrite the volume. The volume is in use by another copy operation such as XRC or TDMF.
- 88 The Peer-to-Peer Remote Copy/FlashCopy session id is not active.
- 89 An establish Peer-to-Peer Remote Copy session command was received, but the device is already associated with a different session.
- 8A A Restore FlashCopy relationship command was received and the volumes were not in a persistent relationship whose background copy had completed.

- 8B Critical Volume mode was specified with an asynchronous Peer-to-Peer Remote Copy operation.
- 8C An Establish Alphabet Peer-to-Peer Remote Copy pair command was received and the pair is not a suspended Alphabet Peer-to-Peer Remote Copy pair.
- 8D An Establish Alphabet Peer-to-Peer Remote Copy pair command was received on a FlashCopy source volume.
- 8E An Establish Alphabet Peer-to-Peer Remote Copy pair command was received on a non-Alphabet Peer-to-Peer Remote Copy primary volume.
- 8F An establish Peer-to-Peer Remote Copy pair command was received and the secondary volume is an Alphabet FlashCopy target volume.
- 90 An establish non-Alphabet Peer-to-Peer Remote Copy pair command was received and the secondary volume is an Alphabet Peer-to-Peer Remote Copy primary volume.
- 91 A Perform Subsystem Function Prepare to Read Lock Data order was received requesting Multi-Path Lock Facility data and the data was not available for presentation because: either the requested Multi-Path Lock Partition is not initialized, or cache is not available for use.
- 92 An Establish Alphabet Peer-to-Peer Remote Copy Pair command was received and the primary volume is a non-Alphabet FlashCopy target volume.
- 93 A cache storage control command cannot be accepted now because a Multi-Path Lock Disconnect operation is in progress for an initialized Multi-Path Lock Partition.
- 94 A command was received specifying that it should be sent to the Peer-to-Peer Remote Copy remote control unit for execution and the volume receiving the command is not a Peer-to-Peer Remote Copy primary.
- 95 A Multi-Path Lock Facility order that can generate an Attention message cannot be accepted because the subsystem message buffer is full.
- 96 An Establish Peer-to-Peer Remote Copy Pair command attempted to perform a transition between asynchronous operation states that is not allowed.
- 97 An DSO Establish FlashCopy order with the restart bit set (byte 1 bit 3) was received and the specified source and/or target volumes are in an existing relationship with some other volume.
- 98 An DSO Establish FlashCopy order with the Incremental Peer-to-Peer Remote Copy bit set (bytes 2-3 bit 10) was received for a volume that is not in the correct Peer-to-Peer Remote Copy state. The Peer-to-Peer Remote Copy Primary volume was not established with bytes 32-33 bit 7 set or the volume is not suspended at the time the DSO Establish FlashCopy order is received.
- 99 A command was received specifying an operation that is not supported on the hardware or software level of ESS.

- > **9A** A DSO Establish FlashCopy order with byte 1 bit 3 set indicating a Restart FlashCopy operation could not be completed due to the inability to write metadata or global status, or some other internal error. The volumes will be placed into a terminate pending state with an indication in NVS to this effect.
- > **9B** A DSO Establish FlashCopy order or Establish Peer-to-Peer Remote Copy Pair order was rejected because the target LUN is reserved to another initiator.
- > **9C** A DSO Establish FlashCopy order cannot complete due to internal hardware conditions or lack of resources to manage the relationship.
- > **9D** A DSO Establish FlashCopy order cannot complete because the maximum number of relationships for the volume has been exceeded. (Data Set level FC only)
- > **9E** PPRC Carousel mode cannot be specified if Alphabet PPRC and FlashCopy Incremental Copy is also specified.
- > **9F** PPRC Carousel mode cannot be specified at the same time as Establish PPRC with no-copy is specified.
- > **A0** A FlashCopy Remote order must specify the remote source volume to be the secondary PPRC volume for the PPRC primary volume this command was addressed to.

| **A1-FF** Reserved

Functional Change

Add a field to Sense which contains a pointer to the bad parameter byte when '04' Invalid Parameter

sense is returned. This would only be used for PSF and DSO commands. - Section 6.3.2.2 Page 208

End of Functional Change

6.3.2.2 Bytes 8-23 for Format 0


The following is a description of the Format 0 sense bytes.

Bytes Definition

- 8** Contains 'Reason Code' in some messages, otherwise zero. See "Byte 1, Bit 1 Invalid Track Format" on page 200 and Message F in the preceding section.
- 9** Contains a bit vector which points to the invalid parameter bits when Sense Format 04 Invalid Parameter. If the error does not involve bits the byte will contain zero. This would be used for PSF and DSO commands only.
- 10-11** Contains a pointer to the invalid parameter byte. The pointer begins at a value of one. If the field is zero, this field is not valid. This would be used for PSF and DSO commands only.
- 12-14** Zeros
- 15-19** Serial Number
Sense bytes 15 through 19 contain the plant of manufacture and serial number information. This data is obtained from the VPD.
- 20-21** Subsystem ID.
- 22-23** Fault Symptom Code (FSC)

For New Requests: Note:
PS5800, PS6000 and DS8000 requests must be created on the new PS61-3, PGP and Requirements database. Refer to the About this Database page under the Help menu for more information.



Need help or instructions? → 

Submitter: William Micka/Tucson/IBM
Date:

Proposed solution: (To be completed by Development)

External design / functional description (including limitations):


Technical design / implementation:

FlashCopy could keep bitmaps for tracking changes to either the source device or the target device of an FlashCopy relationship (on a track granularity). These bitmaps could be used to speed-up subsequent FlashCopy operations for the same relationship (only the changed tracks will be copied). If the customer is performing multiple FlashCopy operations between the same source and the same target in a relatively high frequency (e.g., for data mining) or the amount of changes to either device is small, this proposal will result in a significant performance improvement (significantly less tracks have to be copied).

Additional information by W.Micka. Change recording for FlashCopy and then using the changes to perform an incremental operation is useful only for those subsequent operations that physically move data. For FlashCopy, incremental changes can be used to resync the source or target when the background copy task is specified. For PPRC, the incremental changes will be provided to the primary volume of a pair and be used to resynchronize the copy pair. An illustration of the PPRC operation is provided below.

INCREMENTAL ASYNCHRONOUS PPRC WITH FlashCopy - IRS ACCOUNT

SOLUTION



AT INCREMENTAL BACKUP TIME

1. NEW ESTABLISH FlashCopy Atomic Operation
 - FORCE DESTAGE OF ALL MODIFIED DATA FOR DEVICE 'A'
 - TRANSFER FC BITMAP TO SUSPENDED PPRC PRIMARY
 - FC BITMAP ORed INTO PPRC SUSPENDED BITMAP
 - REMOVE FC RELATIONSHIP A TO B
 - REESTABLISH FC RELATIONSHIP A TO B
 - SELECT NO BKGND COPY
2. NEW ESTABLISH FlashCopy Atomic Operation
 - WITHDRAW AND ESTABLISH FlashCopy C TO D
 - SELECT EITHER NO BKGND COPY OR BKGND COPY
3. NEW ESTABLISH PPRC Operation
 - PPRC RESYNC B TO C WITH AUTOMATIC SUSPENSION

INITIAL SETUP

1. FLASH COPY A TO B
 - SELECT NO BKGND COPY
2. PPRC B TO C
 - DETECT FULL DUPLEX STATE
3. SUSPEND PPRC B TO C
 - SUSPEND BITMAP MODE
4. FlashCopy C TO D
 - SELECT EITHER NO BKGND COPY OR BKGND COPY

W. HIRAKA

IBM CONFIDENTIAL

ASYNCH_FC.PRC

Supported models: ☐ Exx ☒ Fxx ☒ 8xx ☐ Next Gen Minimum LIC level: 2.1.0 Hardware Required: No
MES Creation or Update Required: ☐ Yes

Reason for request:

Marketplace Need / Customer Requirement

Provide an option in FlashCopy whereby only changed tracks are monitored since the initial (or last) Incremental FlashCopy was done. For FlashCopy, resynchronization could occur in either direction, although only in one direction at a time (i.e., original source --> original target OR original target (now the source) --> original source (now the target)). For PPRC, the incremental operation will provide an Asynchronous PPRC solution for the open systems customers which provides consistent data at the remote location.

This feature provides improved refresh time / performance and is currently offered by our competition.

Additional comments:

Related requests:

Additional authors:

Revision history:

- Changed name from FlashCopy; Incremental with resynchronization and PPRC to Remote FlashCopy based on BLM input.
- Primavera
 - Changed submitted from IBM to William Micka/Tucson/IBM
 - Bill Micka updated with the current proposal for Flashcopy+PPRC.
- Updated "Reason for Request" with marketplace need / customer requirement.
- Changed all references of "Instant Image" to "FlashCopy" per PCR P-1136.
- Added design detail and described how incremental change recording can be used for Async PPRC

Prioritized? Yes

(Select "View -> Refresh" or press F9 whenever you change this field.)

Prioritization			
Priority:	High	Prioritized by:	BLM
Rank (within priority):	Medium	Date of last prioritization:	
Status			
Current status:	Reject as of	PM:	
Next review by:	PCR Board on	ESS Rel:	Int GA:
Test Support Link		SW Rel	Int GA:
			SSRdb Link
Review History			

- - PCR Board
- Change Status to Reject
- No plan to pursue this further
- - PCR Board
- Status Plan and Copper Shark target
- - PCR Board
- Change Status to Plan and Copper Shark 3Q02 target
- Comeback
- - Status update (Primavera)
- Change Release to Copper Shark 3Q02 target
- - PCR Board
- Change Status to Plan and Silver Tip+1 3Q02 target
- - PCR Board
- Change Status to Plan and Silver Tip 2Q02 target
- Testing in the 2Q02; PCR 1529 is not a coreq.
- Also applies to F models.
- - PCR Board
- PCR 1529 is a co-requisite with Remote Flash Copy PCR 1059 that will be changed from Plan to Proposal until the complete solution is evaluated.
- Silver Tip DCT has requested that the BLM decide if the Remote FlashCopy solution meets customer requirements
- Comeback
- - PCR Board
- Change Status to Plan and Silvertip 2Q2002 target
- 4 week exposure (2 weeks in code build) and (2 weeks in test) Will work to contain.
- Currently no plan to test this on F Models in 1st or 2nd quarter.
- - PCR Board
- Status Sizing and Silvertip 2Q2002 target
- Gary Albert wants Blacktip to merge the code and test before Silver Tip.
- Comeback
- - Status update (Primavera)
- Will not be done as Blacktip 2; now requested for Silver Tip.
- Status changed from Plan to Sizing; Release changed from Blacktip 2 to Silver Tip.
- - PCR Board
- Change Status to Plan and Blacktip 2 target

- **- PCR Board**
 - Change Status to Review and Blacktip 2 target
 - Comeback
- **- Status update (Primavera)**
 - Reviewed at the PCR Board meeting on 3/8 as PCR 1529.
 - This PCR had been updated with the same information and will be used to track the design/implementation of the Incremental FlashCopy+PPRC proposal.
 - Status Review and Blacktip++1 target
 - Yanes will provide the target code delivery to EVT.
 - Washuk will work with the DCT to establish EVT and test start dates for a November deliverable.
- **- Prioritization**
 - With the inclusion of Incremental FlashCopy and PPRC Asynchronous , this PCR now exists in the top 5. It solves the requirement for number 4 until the real Asynchronous design is completed .
- **- PCR Board**
 - Status Proposal and Blacktip GA target
 - Bill Micka to add technical description and assumptions.
 - Will investigate whether can price this feature separately.
- **- Status update**
 - Change status to Proposal for Blacktip release (per Fail Plan).
- **- Status update**
 - Change status to Plan with Internal GA of Blue Shark BS2 per Fall Plan.
- **- Prioritization**
 - Changed status to Open
 - This PCR was not prioritized by the BLM as one of the top 5 PCRs which should be considered for inclusion in Shark 1.5
 1. P-1050 - Native LSA FlashCopy
 2. P-1011 - PPRC Fibre Channel
 3. P-1043 - SCSI command for FlashCopy and PPRC
 4. P-1044 - Asynchronous PPRC
 5. P-1066 - LSA space reclamation for open systems
 6. P-1065 - File level FlashCopy for open systems (P-1043 is a pre-requisite)
 7. P-1067 - Defragmentation support in FlashCopy
 8. P-1064 - Partition level FlashCopy for open systems
 9. P-1061 - Dataset level FlashCopy for CKD HA LSS
 10. P-1059 - FlashCopy Resynchronization
 11. P-1033 - PPRC read to secondary
 - P-1062 - LSA toleration mode FlashCopy

Category: Copy services

Comments:

Implementation

PCR Evaluated:	Yes	By:	Affected	Status	FCR Number(s)
CKD Attachment Spec:	Yes		Yes	NFCR510 NFCR571 NFCR572	
SCSI Attachment Spec:	Yes		Yes	???	
ESSNet Attachment Spec:	Yes		Yes	Complete	SS1102-01
Functional Spec:	Yes		Yes	Complete	SS1001-08, SS1102-01, FCQR591
Internal Architecture:	No		No		
FCR Process:	Pending				
FCR Comments:					
SS1001-08 Complete, SS1102-01 Complete. Updated to Copper Plan					on FCQR591. Need to check with Micka on S7390 AS, SCSI AS

IBM Confidential

Document last modified on

***Question 3:** What is the scope of the claim?

Reason(s) for above Answer:

Portfolio Need

***Question 1:** What are the portfolio needs in the area of your invention?

Reason(s) for above Answer:

Exploitation & Enforcement

***Question 1:** How easily can the use of the invention by a competitor be detected?

Reason(s) for above Answer:

***Question 2:** How easily can the use of the invention be avoided by a competitor?

Reason(s) for above Answer:

Business Value

***Question 1:** What percentage of the companies producing products in the field of this invention might use this invention?

Reason(s) for above Answer:

***Question 2:** What is the value of this patent to current or anticipated Alliance Activity between IBM and other companies?

Reason(s) for above Answer:

***Question 3:** What is the value of this patent to current or anticipated Technology Transfer Activity between IBM and other companies?

Reason(s) for above Answer:

***Question 4:** Does it result in prestige to IBM?

Reason(s) for above Answer:

Evaluation

Search Information

Search Office Information

Allocation and Additional Billers

Final Decision

Post Disclosure Text & Drawings

(Form Disclosure, Revised 01-Dec-2006)